arXiv:1606.03298v1 [cs.AI] 10 Jun 2016

Structured Factored Inference: A Framework for Automated Reasoning in Probabilistic Programming Languages

Avi Pfeffer
Charles River Analytics
Cambridge, MA 02138
apfeffer@cra.com

Brian Ruttenberg
Charles River Analytics
Cambridge, MA 02138
bruttenberg@cra.com

William Kretschmer*
MIT
Cambridge, MA 02139
kretsch@mit.edu

Abstract

Reasoning on large and complex real–world models is a computationally difficult task, yet one that is required for effective use of many AI applications. A plethora of inference algorithms have been developed that work well on specific models or only on parts of general models. Consequently, a system that can intelligently apply these inference algorithms to different parts of a model for fast reasoning is highly desirable. We introduce a new framework called structured factored inference (SFI) that provides the foundation for such a system. Using models encoded in a probabilistic programming language, SFI provides a sound means to decompose a model into sub–models, apply an inference algorithm to each sub–model, and combine the resulting information to answer a query. Our results show that SFI is nearly as accurate as exact inference yet retains the benefits of approximate inference methods.

1 INTRODUCTION

Probabilistic modeling is at the core of many artificial intelligence (AI) applications. The complexity, richness, and diversity of these models are rapidly growing as AI takes on a larger role in everyday life. As a result, the efficiency of probabilistic inference is critical for effective use of these models. However, despite significant research into efficient algorithms and techniques, probabilistic inference remains a significant bottleneck in many real–world AI applications.

Many different algorithms have been explored to reason on general models. Unfortunately, no single algorithm performs sufficiently on every model, and often there are trade–offs that must be made. For example, sampling methods such as Metropolis-Hastings [1] are often the "go to" algorithms for reasoning on continuous models, but convergence can be painfully slow and they suffer from high variance estimates. Exact methods such as variable elimination [2] work well on discrete problems, but are intractable for all but the simplest models. Recent work on generalized variational inference [3] shows promise, but still requires some hand–tuning to work effectively. Once an algorithm has been found to work well on a specific problem, even slight modifications are no guarantee of continued success; adding a single continuous variable to an otherwise discrete problem can vastly affect the performance of the existing algorithm. Thus, AI developers are faced with selecting and configuring the appropriate algorithm that will work well on their problem, a task that is often more time consuming than constructing the actual model.

One solution to reduce this burden is to develop a method that can automatically select an algorithm that should perform well on one's specific problem. One major impediment to this approach is that the size and complexity of real–world models makes it difficult to determine the best single algorithm. Indeed, different algorithms might be appropriate for different parts of the model. For example, one algorithm might be appropriate for a continuous portion of the model while another is used for a discrete portion; indeed, the Rao–Blackwellization algorithm [4] exploits this fact. As a result, one approach to achieving automated inference is to not just select a single algorithm, but rather a *set* of algorithms to apply to different parts of a model, and combine the results in the appropriate manner. Central to this approach is developing a sound method to decompose a model into manageable sub–models with the appropriate granularity. Such a method would provide a framework to intelligently select algorithms for different parts of a model, and the means to combine results to answer the query.

The emerging field of probabilistic programming (PP) provides the opportunity to support this decomposition framework. PP [5, 6] provides expressive and general purpose languages to encode a probabilistic model as an executable

*Work performed while interning at Charles River Analytics

program. This allows one to leverage the power of programming languages to create rich and complex models, and use built–in inference algorithms that can operate on any model written in the language. More importantly, since the models are encoded as a program, we can use the program structure to understand the properties of a model before inference is even attempted.

We introduce a new PP–based inference framework called structured factored inference (SFI). SFI uses simple PP semantics to identify *decomposition points* within a probabilistic model, creating an abstract hierarchy of sub–models. Each sub–model is independently reduced to a joint distribution over variables relevant to answering the query, using any inference method. Using factors to represent this joint distribution, the results are incorporated into the inference algorithms applied on other sub–models. The SFI framework brings many significant advantages. First, SFI provides the capability to apply *decomposition strategies* to decomposition point so that sub–models can be created with small interfaces (and thus small joint distributions). For example, a strategy could choose to create a sub–model defined by a single decomposition point or combine several decomposition points into a single sub–model. Second, it has the ability to apply *inference strategies* that choose algorithms to "solve" a particular sub–model during the inference process. This means that one need not decide on a single inference algorithm to apply to an entire model.

We show the benefits of SFI on three realistic models using a combination of exact and approximate inference algorithms. Our experiments show that even with simple strategies for decomposition and inference, probabilistic reasoning using SFI achieves performance equal to or better than approximate inference methods and is nearly as accurate as exact inference methods. The SFI framework is extremely general and expandable, providing the opportunity to use more complex decomposition strategies or intelligent inference strategies. The SFI framework has the potential to be the foundation for a general automated inference system.

2 RELATED WORK

Automated algorithm selection has been a long desired goal in computer science, with possibly the first formulation by Rice [7]. As such, it has been applied to a variety of disciplines in the field, such as scientific computing [8], game theory [9], and artificial intelligence problems such as satisfiability [10]. Most efforts, however, focus on methods to analyze and learn how to apply the single best algorithm to solve a problem. For example, Guo [11] uses Bayesian networks to learn and select the best algorithm to solve a problem; neither the problems or algorithms are specific, and can be applied generally to a variety of problems. Our SFI framework is complementary to much of this existing work. SFI can decompose complex models into smaller sub–models that these sophisticated learning algorithms can operate on, potentially providing even better performance than learning on a single model.

Probabilistic inference is unique in some respects as the independence properties of models provides the opportunity to apply many algorithms to different parts of a problem. There has not been significant amount of progress in the probabilistic modeling community, however, to address this or take advantage of it. Our approach is similar in spirit to the current work on black box variational inference [12]. These recent methods have attempted to reduce the programmer burden of configuring and applying variational inference to general probabilistic models, and in a sense are attempting to automatically find the best configuration of the algorithm that produces optimal inference results. While these approaches are promising, they still only consider a single algorithm. The decomposition strategies in SFI also bears similarity to structured variable elimination (SVE) [13]. Like SVE, SFI enjoys the benefits of decomposition in exploiting small interfaces and reusing work. However, SVE applies the same algorithm to each problem, whereas SFI is a general framework for decomposing problems and optimizing each sub–model separately. Finally, our work is similar to decomposition methods that solve maximum a posteriori (MAP) queries presented by Duchi et. al. [14]. This work, however, only applies to specific decompositions of Markov random fields and only applies to MAP problems.

3 PROBABILISTIC PROGRAMMING LANGUAGE

The fundamental concepts in SFI are strongly tied to probabilistic programming languages (PPL), and SFI has been implemented in a publicly available PPL. As such, understanding PPL semantics is critical for understanding SFI. However, since the focus of this paper is introducing the SFI concept, we present a simplified and abstract PPL for purposes of explanation. We call this abstract language SimplePPL.

3.1 SimplePPL Language

The central concept in SimplePPL is a random variable (RV). Intuitively, a RV represents a random process that stochastically produces a value. For simplicity, we use an untyped language, so an RV can produce any value in the value set T, where T is a countable finite set. A program Q has a set of *free* RVs $\mathcal{F}_Q$, and consists of a sequence of definitions of the form $RV = expr$. The set of RVs defined by Q is denoted $\mathcal{RV}_Q$. An RV is *available* if it is either in $\mathcal{F}_Q$ or defined previously in Q. The set of available RVs with respect to an RV r is denoted $\mathcal{A}_r$.

An expression defining an RV r is one of the following:

- A value v
- A primitive defining a probability distribution over values. Examples include $Flip(p)$, which produces true with probability p, and $Uniform(x, y)$.
- $Apply(r_1, \ldots, r_n, f)$, where $r_1, \ldots, r_n$ are available RVs and f is a function $T^n \rightarrow T$.
- $Chain(r_1, f)$, where r_1 is an available RV and f is a function $T \rightarrow \mathcal{Q}$, where $\mathcal{Q}$ is the space of programs such that for each $Q' \in \mathcal{Q}$, $\mathcal{F}_{Q'} \subseteq \mathcal{A}_r$ and the final RV in the program is named "outcome".

3.2 SimplePPL Semantics

Although there are clear semantics for recursive programs in SimplePPL, for simplicity in this paper it will suffice to assume that the expansions of a program Q are non–recursive and thus finite. Under the SimplePPL semantics, each program Q defines a conditional probability distribution $P(\mathcal{RV}_Q|\mathcal{F}_Q)$. This is achieved by defining, for each RV r defined in Q, a conditional distribution $P(r|\mathcal{A}_r)$ and then using the chain rule so that

$$P(\mathcal{RV}_Q|\mathcal{F}_Q) = \prod_{r \in \mathcal{RV}_Q} P(r|\mathcal{A}_r)$$

$P(r|\mathcal{A}_r)$ is defined as follows:

- For r defined by value v, $P(r|\mathcal{A}_r)$ assigns probability 1 to v.
- For r defined by a primitive distribution π, $P(r|\mathcal{A}_r)$ is $P(\pi)$.
- For r defined by $Apply(r_1, \ldots, r_n, f)$, by assumption $\{r_1, \ldots, r_n\} \subseteq \mathcal{A}_r$. Therefore, $P(r|\mathcal{A}_r)$ assigns probability 1 to $f(v_1, \ldots, v_n)$ for any values $v_1, \ldots, v_n$ in the support of $r_1, \ldots, r_n$.
- For r defined by $Chain(r_1, f)$, let v_1 be the value of r_1. Let Q' be the program $f(v_1)$. By induction, Q' defines a distribution $P_{v_1}(\mathcal{RV}_{Q'}|\mathcal{F}_{Q'})$. By definition, $\mathcal{F}_{Q'} = \mathcal{A}_r$ and because "outcome" $\in \mathcal{RV}_{Q'}$, Q' defines a distribution $P_{v_1}(outcome|\mathcal{A}_r)$. As $r_1 \in \mathcal{A}_r$, the distribution $P(r|\mathcal{A}_r)$ is equal to $P_{v_1}(outcome|\mathcal{A}_r)$ for any value $v_1 \in r_1$.

SimplePPL is purely functional, and evidence (conditions or constraints) can be applied as part of a query on a model. While SimplePPL lacks the complexity of many PPLs, it is as expressive as full–fledged PPLs.

4 STRUCTURED FACTORED INFERENCE

4.1 Overview

There is a simple intuition behind SFI: If a model can be broken down into smaller sub–models (i.e., programs) that can solved independently (i.e., marginalizing out non–relevant variables), then different algorithms can be applied to different parts of a model. Combined with methods for intelligent inference algorithm selection, this framework could then lead to improved inference on a wide variety of problems. SFI is fundamentally a framework for applying two types of strategies: A decomposition strategy that divides a model into smaller sub–models, and an inference strategy that appropriately applies an inference algorithm to each sub–model. SFI uses factors to combine information from solved sub–models to answer queries.

```
a = Flip(0.6)
b = Chain(a, f)
c = Chain(a, f)

f(true) = {
  x1 = Flip(0.9)
  y1 = Flip(0.8)
  z1 = Apply(x1, y1, (b1, b2) => b1&&b2)
  x2 = Flip(0.7)
  y2 = Flip(0.8)
  z2 = Apply(x2, y2, (b1, b2) => b1&&b2)
  outcome = Apply(z1, z2, (b1, b2) => b1||b2)
}
f(false) = {
  x1 = Flip(0.1)
  y1 = Flip(0.8)
  z1 = Apply(x1, y1, (b1, b2) => b1&&b2)
  x2 = Flip(0.2)
  y2 = Flip(0.8)
  z2 = Apply(x2, y2, (b1, b2) => b1&&b2)
  outcome = Apply(z1, z2, (b1, b2) => b1||b2)
}
```

Figure 1: SimplePPL code for a small model. The "=>" denotes an anonymous function in SimplePPL.

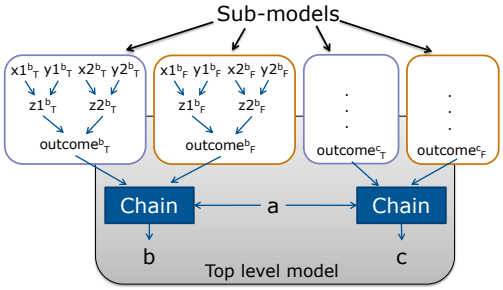

Figure 2: The model presented in Fig. 1

As an example, consider the following model written in SimplePPL, shown in Figs. 1 and 2. We have three RVs defined in Q, a, b, and c. RVs b and c generate a value using the Q' that is generated by $f(a)$, where $outcome^b_T$ refers to the outcome for RV b when a is true, and so forth. Each Chain generates a program Q' for both true and false conditions. With the exception of "outcome", the RVs defined by Q' are not directly needed to reason about a, b, and c. That is, all of the RVs defined in Q' except "outcome" and $\mathcal{F}_{Q'}$ can be marginalized out of Q'. A joint distribution over "outcome" $\cup$ $\mathcal{F}_{Q'}$ is all that is needed to reason at the top–level program Q. Since this marginalization is self–contained, any inference algorithm that can compute a

joint distribution can be applied to each sub–model (shown in the boxes in Fig. 2). This joint distribution represented as a factor; these factors are then "rolled up" and used by another algorithm to answer queries on the program Q. This is the core operation of SFI: Given a sub–model, use an algorithm to marginalize away internal variables and return a joint distribution over "outcome" and the free variables, and repeat the process until the query is answered.

4.2 Model Decomposition

The following discussion of SFI is cast in the context of a SimplePPL program. However, the method is applicable to graphical models in general.

The key operation of SFI is model decomposition. This operation decomposes a model into semantically meaningful sub–models (i.e., programs) that can be reduced to a joint distribution over relevant variables. First, we define two key concepts: *uses* and *external*. A RV r *uses* a RV x if:

$$x \in \mathcal{A}_r \wedge P(r|\mathcal{A}_r) \neq P(r|\mathcal{A}_r \setminus x)$$

We denote the set of variables $\mathcal{U}_r$ for a variable r as the set of all variables r uses, either directly or recursively, plus r. This definition of uses can be difficult to verify in a program based on the semantics of SimplePPL. However, in our implementation of SimplePPL, we have a syntactic (but stronger) condition for r using x based on x appearing in the expression for r, or in any expression for a variable used by r. Such a condition is necessary for uses and thus still guarantees SFI's soundness.

We denote that a variable x is *external* to r if:

$$x \in \mathcal{U}_r \wedge \exists\ y \in \mathcal{RV}_Q \setminus \mathcal{U}_r \,|\, x \in \mathcal{U}_y$$

That is, an external variable to r is used in the generative process of a variable that r does not use. We denote the set of variables external to r as $\mathcal{E}^r$.

A decomposition of the model with respect to a RV $d \in \mathcal{RV}_Q$ is an operation that partitions $\mathcal{RV}_Q$ into two disjoint sets of variables, $\mathcal{RV}_Q^d$ and $\mathcal{RV}_Q^{\overline{d}}$. The RV d is called a decomposition point. We define $\mathcal{RV}_Q^d$ and $\mathcal{RV}_Q^{\overline{d}}$ as:

$$\begin{aligned} \mathcal{RV}_Q^d &= \mathcal{U}_d - \mathcal{E}^d \\ \mathcal{RV}_Q^{\overline{d}} &= \mathcal{RV}_Q - \mathcal{RV}_Q^d \end{aligned} \tag{1}$$

In other words, $\mathcal{RV}_Q^d$ is the set of variables exclusively used in the generation of d (i.e., no external uses), and $\mathcal{RV}_Q^{\overline{d}}$ is all remaining variables in Q.

As an example, consider the program in Fig. 1. Since d can be any variable, let us choose $outcome_T^b$ as the decomposition point. In this example, $\mathcal{RV}_Q^d$ is the set of variables that $outcome_T^b$ exclusively uses, so it would be $\{outcome_T^b, x1_T^b, y1_T^b, z1_T^b, x2_T^b, y2_T^b, z2_T^b\}$ (all the variables in the left–most box of Fig. 2). $\mathcal{RV}_Q^{\overline{d}}$ would include all other variables in the model.

4.3 Factored Representation

Using factors in SFI has several advantages. First, it provides an interface to communicate the joint distribution of a sub–model to other parts of the model. Second, factors make SFI algorithm–agnostic; any algorithm that can compute a joint distribution and return a factor can be used. For example, sampling algorithms that can post–process a joint distribution into a factor can also be used.

4.3.1 Factor Creation in SimplePPL

Once the variables in Q have been split into two sets via a decomposition point, we convert the decomposition to a factored representation. Each variable $r \in RV_Q$ can be converted to a set of factors Ψ_r that describe the generative semantics of the variable. For variables defined as values or primitives, we create a factor over the support of the variable using the probability distribution defining the variable. For variables defined by $r = Apply(r_1, \ldots, r_n, f)$, we create a single factor over $\{r, r_1, \ldots, r_n\}$ whose value is 1 when $r = f(r_1, \ldots, r_n)$ and 0 otherwise.

Finally, for variables defined by $r = Chain(r_1, f)$, we create a set of factors that represents the joint distribution of $\{r, r_1, outcome_1 \ldots, outcome_n\}$. Because representing the joint distribution of $\{r, r_1, outcome_1 \ldots, outcome_n\}$ could be prohibitively large, we decompose this joint distribution to keep the sum–product operations tractable. For each value v of r_1, we create a factor ψ_v over $\{r, r_1, outcome_v\}$, where $outcome_v$ is the "outcome" variable generated from applying $f(v)$. We then generate probabilities for the factor in the following manner:

- For each $v' \in r_1$, if $v' \neq v$, then the probability is 1. This is a "dont care" case.
- For each $v' \in r_1$, $w \in r$, and $x \in outcome_v$, if $v' = v$ and $w = x$, then the probability is 1.
- Otherwise, the probability is 0.

We also create a selector factor over $\{r, r_1\}$ that selects a factor over $outcome_v$ for the appropriate value of $f(v)$.

4.3.2 SFI with Factors

We denote the set of factors created from a program Q as Ψ, and each factor $\psi \in \Psi$ is defined over a set of variables $x_\psi \in \mathcal{RV}_Q$. Given Ψ, the probability distribution of a RV r in the program, $P(r)$, is formulated as:

$$P(r) = \frac{1}{Z} \sum_{x \in \mathcal{RV}_Q \setminus r} \prod_{\psi \in \Psi} \psi(x_\psi) \tag{2}$$

where Z is the normalizing constant.

As d divides $\mathcal{RV}_Q$ into two sets, it naturally divides Ψ into two sets, Ψ_d and $\Psi_{\overline{d}}$, as explained in Sec.r̃effactorCreation. As such, with $y_\psi \in \mathcal{RV}_Q$, we can rewrite Eqn. 2 as:

$$P(r) = \frac{1}{Z} \sum_{x \in \mathcal{RV}_Q^{\overline{d}} \setminus r} \sum_{y \in \mathcal{RV}_Q^{d} \setminus r} \prod_{\psi \in \Psi_{\overline{d}}} \psi(x_\psi) \prod_{\psi \in \Psi_d} \psi(y_\psi) \tag{3}$$

Note that even though the variables in $\mathcal{RV}_Q^d$ and $\mathcal{RV}_Q^{\overline{d}}$ are disjoint, the variables used in the sets of factors Ψ_d and $\Psi_{\overline{d}}$ are *not* disjoint. From the definition of the sets in Eqn. 1, the only variables that are shared between Ψ_d and $\Psi_{\overline{d}}$ can be $\mathcal{E}^d$, the set of variables external to d, and d itself. As such, we can move the summation over $\mathcal{RV}_Q^d$ in Eqn. 3 inwards and the summation over d to the outer summation, so that we get:

$$\begin{aligned} P(r) &= \frac{1}{Z} \sum_{x \in \{\mathcal{RV}_Q^{\overline{d}} \cup d\} \setminus r} \prod_{\psi \in \Psi_{\overline{d}}} \psi(x_\psi) \sum_{y \in \{\mathcal{RV}_Q^{d} \setminus d\} \setminus r} \prod_{\psi \in \Psi_d} \psi(y_\psi) \\ &= \frac{1}{Z} \sum_{x \in \{\mathcal{RV}_Q^{\overline{d}} \cup d\} \setminus r} \prod_{\psi \in \Psi_{\overline{d}}} \psi(x_\psi)\, \psi^{\mathcal{E}^d} \end{aligned}$$

where $\psi^{\mathcal{E}^d}$ is a joint factor over d and the external variables defined with respect to d. Again looking at Fig. 1 as an example, with $outcome_T^b$ as the decomposition point, we perform the summation over $\{x1_T^b, y1_T^b, z1_T^b, x2_T^b, y2_T^b, z2_T^b\}$ and are left with a factor over only $outcome_T^b$. This factor can them be multiplied with the remaining factors in the program and $outcome_T^b$ can be summed out.

In this formulation, a decomposition point d implies a structured process to compute $P(r)$ from a set of factors defined on a model: First, compute a joint distribution with respect to the decomposition point, then compute $P(r)$ using the joint distribution and the remaining factors. Computing the joint distribution over the external variables can be accomplished by any algorithm, as can the computation of $P(r)$ once the joint distribution is computed.

So far, we have only mentioned a single decomposition point in a model. However, multiple decomposition points can be defined on a model. In Fig. 2, for example, there are four natural decomposition points ($outcome_T^b, outcome_F^b, outcome_T^c, outcome_F^c$) that can be marginalized independently (shown as the boxes in Fig. 2). Eqn. 3 can be reformulated for multiple points as:

$$P(r) = \frac{1}{Z} \sum_{x \in \{\mathcal{RV}_Q^{\overline{D}} \cup D\} \setminus r} \prod_{\psi \in \Psi_{\overline{D}}} \psi(x_\psi) \prod_{k=1}^{n} \psi^{\mathcal{E}^{d_k}} \tag{4}$$

where there are n decomposition points, and $\mathcal{RV}_Q^{\overline{D}}$ is the intersection of $\mathcal{RV}_Q^{\overline{d_k}}, k = 1, \ldots n$. Decomposition points can be nested inside other decomposition points, allowing inference to proceed in any hierarchical structure implied by the model.

Algorithm 1 Overview of the SFI algorithm

```
function DECOMPOSE(program Q, variables E, dStrategy DS, iStrategy IS)
    Ψ ← ∅
    for c ∈ Q_chain, v ∈ r_c do
        Q' ← f_c(v)
        E_Q' ← E_Q'^{Q'.outcome} ∪ Q'.outcome
        Ψ^{E_Q'} ← DS(Q', E_Q', DS, IS)
        Ψ ← Ψ ∪ Ψ^{E_Q'}
    end for
    ψ^E ← IS(Ψ ∪ Ψ_D̄, E)
    return ψ^E
end function
function SFI(program Q, query q, dStrategy DS, iStrategy IS)
    ψ_q ← Decompose(Q, q, DS, IS)
    return Normalize(ψ_q)
end function
```

In principle, any RV could potentially be a decomposition point. However, we would like to choose a decomposition point d that leads to a small joint factor $\psi^{\mathcal{E}^d}$ and eliminates as many variables in $\mathcal{RV}_Q$ as possible. Chains present a natural decomposition point, which have the benefit of being automatically derived from the program and don't need to be specified by the programmer. When we apply the Chain function f to a parent value v, $f(v)$ is a program that defines a sequence of RVs, ending in a definition of a variable named "outcome". By the semantics of Chain, only the outcome RV can be used anywhere else in the program. For each Chain defined in Q, we create a decomposition point at $outcome_v$ for every value v in the support of r_1. This also implies that $\mathcal{E}^d = \mathcal{F}_{Q'}$ for a decomposition point. Thus, we know that the joint factor created at each decomposition point will only be over each "outcome" variable and free variables defined in the program Q' generated from f.

5 USING SFI

5.1 SFI Operation

Algorithm 1 outlines inference in SFI. To query for the distribution over an RV q, a user calls the SFI function with the program Q (written in SimplePPL), q, a decomposition strategy DS, and an inference strategy IS. DS and IS are functions that guide the decomposition and inference of the model, and are explained in more detail below. The SFI function calls the $Decompose$ function, and the resulting factor over q is normalized to compute $P(q)$.

The $Decompose$ function visits each decomposition point in Q, applies DS to the sub–model (i.e., program) defined

Algorithm 2 A recursive decomposition strategy

function RECURSIVEDECOMPOSITION(program Q, variables $\mathcal{E}$, iStrategy IS)
  return $Decompose(Q, \mathcal{E}, this, IS)$
end function

by each point, and marginalizes out the internal variables using IS. On lines 3, SFI iterates over all Chains defined in Q and each value v of the parent variable r_c. On each iteration, it generates Q', the program created by applying the function f_c to a value v (line 4). Next, it creates the set of relevant variables to program Q' as the external variables in Q' and the "outcome" variable (line 5). It then invokes DS on the new program, which returns a set of factors that is added to the current set for Q (lines 6 and 7). Note that a decomposition may also be recursive, as described below.

Once all decomposition points have been visited, the set of factors not generated from a decomposition point ($\Psi_{\overline{D}}$, via the factor generation described in Sec. 4.3.1) is added to Ψ and IS is applied which returns a factor $\psi^{\mathcal{E}}$ over the variables in $\mathcal{E}$ (lines 9 and 10). Much of the work of the SFI framework is performed by the decomposition and inference strategies, so we explain these in detail below.

5.2 Strategies for Decomposition

A decomposition strategy DS is a method that defines how a program should be decomposed. It is a function that receives a program Q' and set of relevant variables $\mathcal{E}_{Q'}$, and returns a set of factors $\Psi^{\mathcal{E}_{Q'}}$ over *at least* $\mathcal{E}_{Q'}$. The simplest DS is what we call "raising": For a point d, return the set of factors over all variables defined in Q'. This strategy performs no inference, and as a result, all of the factors from $\mathcal{RV}_{Q'}$ are "rolled up" to the top–level. If each d is raised, we get a "flat" strategy. This is how typical inference works; factors are created for all variables and all non–query variables are marginalized out in a flat operation.

To take advantage of different inference algorithms, it is clearly beneficial to have a DS that actually reduces the number of variables in the returned factor set $\Psi^{\mathcal{E}_{Q'}}$. As such, we define a recursive strategy as one that will recursively apply the $Decompose$ function until no more decomposition points are found, shown in Alg. 2.

Here, each decomposition point in a model is recursively visited in a depth–first traversal. Once a program is reached with no decompositions, IS is applied to the factors in the program, and a joint distribution over the external variables and outcome is returned, and the process is repeated. This is referred to as hierarchical inference in SFI.

More complex and sophisticated strategies can also be applied. For example, a strategy could decompose only if $\mathcal{E}$ is at most n variables (n would have to be specified at compile time). If the number of external variables is greater than n, then the function returns all of the factors defined for the program without running any inference strategy. Otherwise, it calls $Decompose$ again to continue the recursive decomposition.

5.3 Strategies for Inference

A strategy for inference applies an inference algorithm to a set of factors defined by program Q and returns a joint distribution over $\mathcal{E}$, the set of external variables in the factors. While SFI uses factors communicate the joint distribution to other programs, there is no restriction that an algorithm operate on factors. As long as the algorithm can ingest factors from other decompositions and output a joint factor over $\mathcal{E}$, then any algorithm can be used.

SimplePPL's implementation of SFI uses factor–based algorithms. There are three algorithms available: Variable elimination (VE) [2], belief propagation (BP) [15] and Gibbs sampling (GS) [16]. VE and BP are standard implementations of these algorithms on factors, and as such we do not provide any details. GS is implemented on a set of factors, but integrating it into SFI is not trivial. Much of the effort is due to the determinism frequently found in PPLs. Our implementation uses automated blocking schemes to ensure proper convergence of the Markov chain. Details on GS can be found in the supplementary material.

5.3.1 Choosing an Algorithm

SFI provides the opportunity to develop schemes that dynamically select the best inference algorithm for a decomposition point, serving as the foundation for an automated inference framework. At the application of the inference strategy, there is opportunity to analyze and estimate the complexity of various algorithms applied to the factors, and choose the one with the smallest cost (e.g., speed or accuracy). For example, methods that estimate the running time of various inference algorithms on a model [17] can be encoded into an inference strategy, and the algorithm with the lowest estimated running time can be chosen.

We created a simple heuristic to choose an inference algorithm, but yet still demonstrates the potential of the approach. As VE is an exact algorithm, it is always preferred over other algorithms if it is not too expensive, but unfortunately is impractical on most problems. We therefore have a heuristic to use VE on a set of factors. We first compute an elimination order, O, to marginalize to the external variables. The cost of eliminating a single variable is the increase in cost between the new factor involving the variable and the existing factors, and the cost of VE is the maximum over all variables, using O. If the cost is less than some threshold we use VE, otherwise, BP or GS.

To choose between BP and GS, we also use another heuristic. As the degree of determinism in a model strongly cor-

relates with the convergence rate of BP [18], we use the amount of determinism in the model as a choice between using BP or GS. We mark a variable as deterministic if, when using GS, we must create a block of more than one variable. If the fraction of deterministic variables (as compared to all variables) in the model is greater than a threshold, then we invoke BP and otherwise GS. While these strategies are heuristics, they do demonstrate the proof of concept for automated inference, and the results presented in the next section show that they are effective.

6 EXPERIMENTS

We tested SFI using three models. First, we encoded a version of the QMR medical diagnosis model [19] in SimplePPL. Like the standard QMR model, this one is a Bayesian network of causal diseases associated with observable symptoms. However, this model inserts a layer of intermediate diseases between the causal and symptom layer. Thus the intermediate diseases are conditioned on the causal diseases, and the symptoms conditioned on the intermediate diseases. The number of diseases and the number of parents per symptom are varied during testing, and the network is constructed by randomly connecting symptoms to intermediate diseases then subsequently the intermediate diseases to random causal diseases. In each test, a random number of symptoms and causal diseases are observed as evidence.

The second model is a mixed directed–undirected model. The undirected portion of the model is an Ising model [20], where each Boolean variable v in an $n \times n$ grid has a potential to its four vertical and horizontal neighbors. The prior over each variable, however, is a modeled as a small Bayesian network conditioned on a causal variable c_n. Thus this model can be viewed as n^2 Bayesian networks that are joined together in a top–level Ising model. The grid size is varied during testing, but for each test a random 20% of the c_n variables are observed as either true or false.

The last model we used is a simplified version of the Bayesian seismic monitoring system presented in [21]. This model is designed to detect and localize seismic events on a simulated two–dimensional representation of earth (with semi–realistic physics). The model consists of a set of monitoring stations at different locations that detect a variety of seismic signals over time. Based on a generative process for both true seismic events and false detections, the model is designed to infer the actual number of seismic events using measurement data (i.e., observations) from each of the detection stations. Continuous variables from the original model were discretized for factored inference; for most tests each distribution was discretized to five bins, but some tests varied the number of bins. We used 10 detection stations with one true event, and varied the total number of false detections from zero to 27. Observation data was generated from a third–party simulation of the seismic generative process. Note that as the number of false detections and discrete bins increases, the model quickly scales up in the number and size of the factors. On some tests, we were unable to attain ground truth since no exact algorithms could complete with available memory.

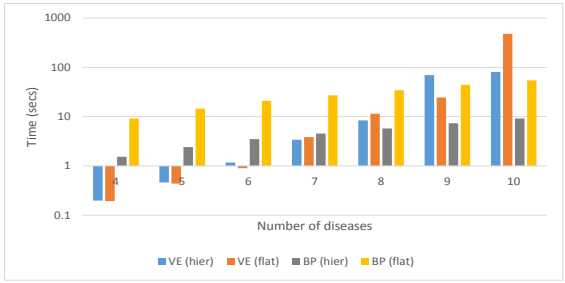

(a)

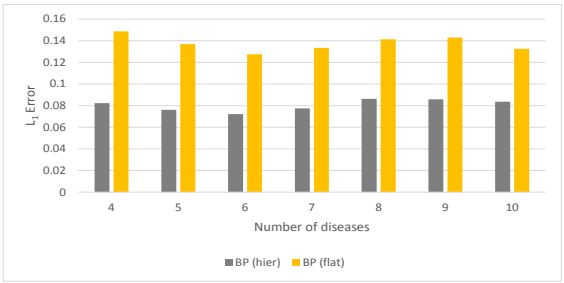

(b)

Figure 3: Comparing (a) running time of flat and hierarchical strategies on VE and BP and (b) accuracy of BP as compared to VE.

All of these models are well suited for SFI. First, they contain a significant amount of structure that can be used for decomposition; the diseases and symptoms in the QMR model represent a series of Chain RVs, each Boolean variable in the Ising model is a Chain that uses each c_n as a parent, and each seismic station's observations are Chain RVs. Second, they are fairly common models used in practice and are realistic. All queries for testing were posterior distribution queries over a random subset of diseases (QMR model), Boolean variables (Ising model), or the number of true seismic events.

6.1 Results

6.1.1 QMR and Mixed Model Results

We first tested how different decomposition strategies affect inference. We used two strategies, flat and hierarchical, on the QMR model using BP and VE. For the hierarchical version of BP, each instance of BP ran for 10 iterations, whereas on the flat version BP was run once for 100 iterations. The results of the running time and accuracy (of BP) are shown in Figs. 3(a) and 3(b), respectively.

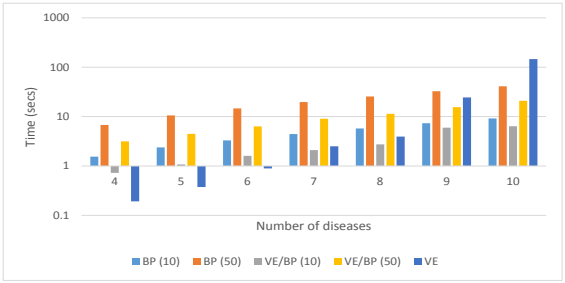

(a)

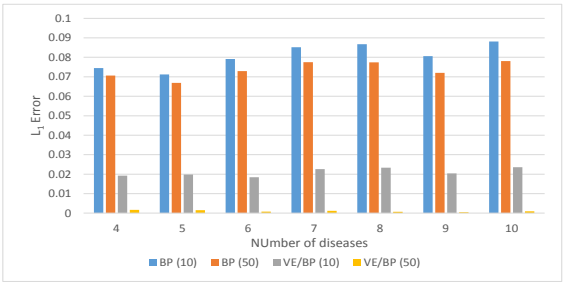

(b)

Figure 4: Comparing the hybrid inference strategy on the QMR model. (a) Running time, and (b) accuracy as compared to VE.

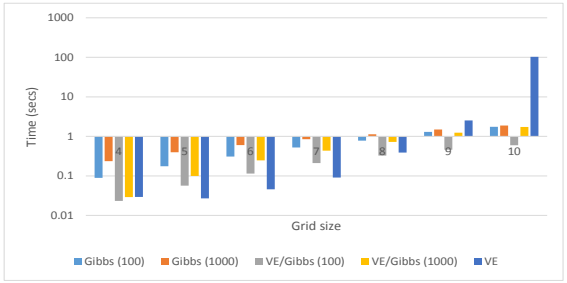

(a)

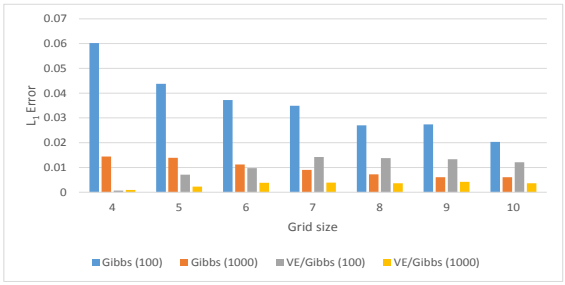

(b)

Figure 5: Comparing the hybrid inference strategy on the mixed model. (a) Running time, and (b) accuracy as compared to VE.

On VE, the results show that the hierarchical strategy generally is faster than the flat one. Mathematically both strategies are performing the same operations. However, the hierarchical strategy imposes a partial elimination order; an elimination order is found for each instance of VE, but the order that the decomposition points are visited is fixed by the strategy. The flat strategy uses a heuristic to find the best elimination order given all the factors in the model. From these results, it appears that the structure imposed by programmer (i.e., by using Chains) finds a better elimination order than the heuristics used to solve this NP–hard problem [22]. These results are consistent with previous work on structured VE [13].

For BP, the hierarchical strategy consistently runs faster than the flat strategy. This comparison is not exact though as it is hard to determine how many iterations in a flat strategy "equals" hierarchical iterations. However, looking at the accuracy of the methods in Fig. 3(b), the hierarchical method is consistently more accurate. Thus, even if the flat BP is run for more iterations to improve its accuracy (assuming it has not converged), it is already dominated in running time and accuracy by the hierarchical method.

Next, we applied our hybrid strategy (with hierarchical decomposition strategies) to determine if we can improve the speed and/or accuracy using multiple algorithms as compared to a single algorithm. The results on the QMR model for running time and accuracy are shown in Figs. 4(a) and 4(b), respectively. The inference strategy only chose to run VE and BP during inference, so we only compare to those algorithms.

For running time, VE remains competitive until the number of diseases reaches about 9. Again, comparing BP to the combined VE/BP method directly is difficult. However, combined with the accuracy results in Fig. 4(b), we can analyze the relationship between running time and accuracy for all the methods. Both BP–10 and BP–50 have comparable accuracy. The hybrid methods, however, are much more accurate. The hybrid VE/BP–10 method is nearly 4 times more accurate than BP–10/BP–50 and VE/BP–50 approaches nearly zero error. The running time for the hybrid methods are both faster than their "respective" BP version (i.e., comparing BP–10 to the hybrid VE/BP–10 test). While VE/BP–50 has a longer running time than the single BP–10 iteration strategy, it is nearly as accurate as VE with significantly less running time.

The results on the mixed model are also shown in Figs. 5(a) and 5(b). The strategy only chose to run VE and GS during inference, so we only compare to those algorithms. Similar to the QMR model, VE has the best performance until the model becomes large, at which point the hybrid strategy has the best running time. Looking at Fig. 5(b), the accuracy of the hybrid methods (GS–100 and GS–1000 iterations) is better than the single GS approaches. Overall, the hybrid approach dominates the GS approach in terms of accuracy and time. However, as more GS iterations are performed, this performance gap will decrease, as the running times of both the single and hybrid approaches are dominated by applying GS to the undirected portion of the model.

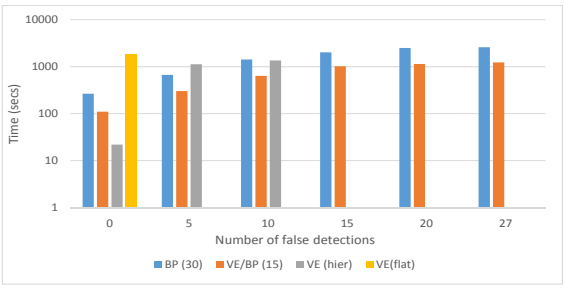

Figure 6: Comparing the running time as false detections increase. VE did not complete on several tests.

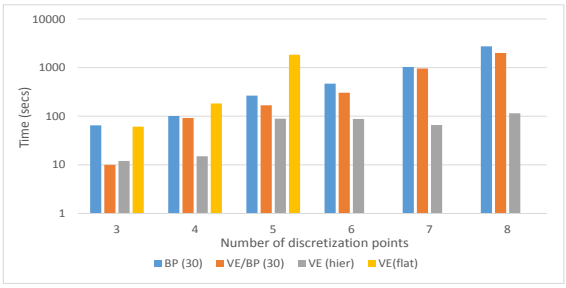

(a)

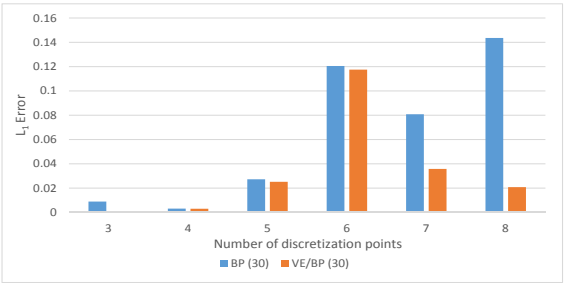

(b)

Figure 7: Comparing the (a) running time, and (b) accuracy as compared to VE as the discretization of the model increases. VE did not complete on several tests.

6.1.2 Seismic Monitoring Results

On the seismic monitoring model, we ran two experiments where we either varied the number of false detections (Fig. 6) or discretization in the models (Fig. 7). Fig. 6 shows the running time of our hybrid VE/BP–15 strategy, hierarchical VE, and flat versions of VE and BP–30 as the number of false detections in the seismic model increases. That is, each test increased the noise in the model, making it harder to detect the true seismic events and significantly increasing the complexity of inference. On the hybrid strategy tests, the SFI framework ran VE for most of the sub–models with the exception of BP on a few sub–models with large tree–widths. Because the two versions of VE ran out of memory on several tests, we are unable to determine ground truth values. The difference in accuracy between VE/BP–15 and BP–30, however, is negligible. We can see that the hybrid strategy has a smaller running time than all the other strategies. As before, it is hard to determine "equivalence" between VE/BP–15 and BP–30, but even if 30 iterations is enough to reach convergence, the hybrid strategy still dominates in terms of time.

Figs. 7(a) and 7(b) compare strategies as the number of discretization points in the model increases. This has the effect of increasing the size of the factors while keeping the number of factors in the model constant (and assumed no false detections). In Fig. 7(a), hierarchical VE actually has the fastest running time. The hybrid strategy is faster than BP–30 and flat VE, but hierarchical VE clearly dominates the hybrid strategy. Again, this demonstrates that programmer imposed elimination orderings can be much better than heuristic approaches. Clearly, however, this shows that our heuristic to choose an algorithm is not sophisticated enough to understand that VE is preferred on this model. A strategy that uses more than just the increase in factor size to select an algorithm is required in this situation.

Finally, since structured VE completed on all tests, we show the accuracy of the VE/BP–30 and BP–30 methods in Fig. 7(b). As can be seen, the hybrid method that uses a combination of VE and BP is more accurate than flat BP. In addition, as shown in Fig. 7(a), the VE/BP–30 method dominates BP–30 in terms of running time.

7 CONCLUSION

In this work, we described a new framework for inference in probabilistic modeling called Structured Factored Inference. Leveraging the capabilities of probabilistic programming, we have shown a semantically sound method to decompose a model into smaller sub–models, and detailed how the application of strategies can guide the inference process. Using simple heuristics to analyze the complexity of sub–models, we demonstrated that SFI can be used to implement a basic automated inference scheme that reasons faster than approximate inference and is nearly as accurate as exact inference methods.

This work serves as a starting point for a more robust automated inference framework, but more analysis is still required. First, there needs to be more theoretical analysis on the criteria for declaring that a sub–model is "solved." That is, how many iterations of GS or BP should be invoked on each sub–model? Most likely answering this question touches upon issues of algorithm convergence, but its impact in the SFI framework is an open research problem.

Second, as shown by the seismic monitoring model, more intelligent algorithm selection methods need to be developed. Recent work provides a starting point for this ([23, 17]), but new methods that leverage the analyzability of PP can make the estimation of complexity more accurate [24]. Finally, new decomposition points would also need to be developed to enable more sophisticated model

decomposition. Chain decomposition is effective, but user–defined or object–oriented decomposition points may be more effective at decomposing a model into sub–models that facilitate faster inference. Our hope is that the SFI framework will be the catalyst for future research in these areas.

Acknowledgements

This work was supported by DARPA contract FA8750-14-C-0011.

References

[1] S. Chib and E. Greenberg, "Understanding the metropolis-hastings algorithm," *The american statistician*, vol. 49, no. 4, pp. 327–335, 1995.

[2] D. Koller and N. Friedman, *Probabilistic graphical models: principles and techniques*. MIT press, 2009.

[3] M. J. Wainwright and M. I. Jordan, "Graphical models, exponential families, and variational inference," *Foundations and Trends® in Machine Learning*, vol. 1, no. 1-2, pp. 1–305, 2008.

[4] A. Doucet, N. De Freitas, K. Murphy, and S. Russell, "Rao-blackwellised particle filtering for dynamic bayesian networks," in *Proceedings of the Sixteenth conference on Uncertainty in artificial intelligence*. Morgan Kaufmann Publishers Inc., 2000, pp. 176–183.

[5] D. Koller, D. McAllester, and A. Pfeffer, "Effective bayesian inference for stochastic programs," in *AAAI/IAAI*, 1997, pp. 740–747.

[6] N. D. Goodman, "The principles and practice of probabilistic programming," in *Proceedings of the 40th annual ACM SIGPLAN-SIGACT symposium on Principles of programming languages*. ACM, 2013, pp. 399–402.

[7] J. R. Rice, "The algorithm selection problem," *Advances in Computers*, vol. 15, pp. 65–118, 1976.

[8] E. N. Houstis, A. C. Catlin, J. R. Rice, V. S. Verykios, N. Ramakrishnan, and C. E. Houstis, "Pythia-ii: a knowledge/database system for managing performance data and recommending scientific software," *ACM Transactions on Mathematical Software (TOMS)*, vol. 26, no. 2, pp. 227–253, 2000.

[9] A. Guerri and M. Milano, "Learning techniques for automatic algorithm portfolio selection," in *ECAI*, vol. 16, 2004, p. 475.

[10] L. Xu, F. Hutter, H. H. Hoos, and K. Leyton-Brown, "Satzilla: portfolio-based algorithm selection for sat," *Journal of Artificial Intelligence Research*, pp. 565–606, 2008.

[11] H. Guo, "A bayesian approach for automatic algorithm selection," in *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI03), Workshop on AI and Autonomic Computing, Acapulco, Mexico*, 2003, pp. 1–5.

[12] R. Ranganath, S. Gerrish, and D. M. Blei, "Black box variational inference," *arXiv preprint arXiv:1401.0118*, 2013.

[13] D. Koller and A. Pfeffer, "Object-oriented bayesian networks," in *Proceedings of the Thirteenth conference on Uncertainty in artificial intelligence*. Morgan Kaufmann Publishers Inc., 1997, pp. 302–313.

[14] J. Duchi, D. Tarlow, G. Elidan, and D. Koller, "Using combinatorial optimization within max-product belief propagation," in *Advances in Neural Information Processing Systems 19: Proceedings of the 2006 Conference*, vol. 19. MIT Press, 2007, p. 369.

[15] J. S. Yedidia, W. T. Freeman, and Y. Weiss, "Understanding belief propagation and its generalizations," *Exploring artificial intelligence in the new millennium*, vol. 8, pp. 236–239, 2003.

[16] S. Geman and D. Geman, "Stochastic relaxation, gibbs distributions, and the bayesian restoration of images," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, no. 6, pp. 721–741, 1984.

[17] L. H. Lelis, L. Otten, and R. Dechter, "Predicting the size of depth-first branch and bound search trees," in *Proceedings of the Twenty-Third international joint conference on Artificial Intelligence*. AAAI Press, 2013, pp. 594–600.

[18] A. T. Ihler, J. Iii, and A. S. Willsky, "Loopy belief propagation: Convergence and effects of message errors," in *Journal of Machine Learning Research*, 2005, pp. 905–936.

[19] T. S. Jaakkola and M. I. Jordan, "Variational probabilistic inference and the qmr-dt network," *Journal of artificial intelligence research*, pp. 291–322, 1999.

[20] R. J. Glauber, "Time-dependent statistics of the ising model," *Journal of mathematical physics*, vol. 4, no. 2, pp. 294–307, 1963.

[21] N. S. Arora, S. Russell, and E. Sudderth, "Netvisa: Network processing vertically integrated seismic analysis," *Bulletin of the Seismological Society of America*, vol. 103, no. 2A, pp. 709–729, 2013.

[22] J. Pearl, *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann, 2014.

[23] N. Flerova, R. Marinescu, and R. Dechter, "Preliminary empirical evaluation of anytime weighted and/or best-first search for map," in *Proceedings of 4th NIPS workshop on Discrete Optimization in Machine Learning*. Citeseer, 2012.

[24] C.-K. Hur, A. V. Nori, S. K. Rajamani, and S. Samuel, "Slicing probabilistic programs," in *ACM SIGPLAN Notices*, vol. 49, no. 6. ACM, 2014, pp. 133–144.